Wiring design helps core memory work at rapid cycle time

An unusual layout in a 2½-dimensional organization gives a ferrite-core computer memory a 500-nanosecond cycle, and keeps cost down

By Alexander Elovic

Electronic Components Division, Burroughs Corp., Plainfield, N.J.

Can a ferrite-core memory's cycle time be pushed well below a microsecond without boosting the cost an unreasonable amount? This question confronted engineers at the Burroughs Corp. Their answer was a new design that allows a memory of 8,192 twenty-bit words to read and write in only 500 nanoseconds. The design clips more than 100 nanoseconds from previous speeds for memories in this size range.

One design considered by the engineers was the conventional 3-dimensional 4-wire design. Many large memories have this organization [see "Memory organizations," p. 85]. But certain problems cropped up. They grew more serious when engineers worked with a cycle time of less than one microsecond. The most serious difficulties, which ruled out further consideration of the 3-D organization, were the distributed capacitance between the drive and inhibit wires, and the additional time that the inhibit pulse imposed on the cycle.

Another approach considered by Burroughs was the 2½-D configuration. This design eliminated the inhibit difficulties, but other problems inherent in all high-speed memories appeared. First, the drive wires can resemble transmission lines if they are physically or electrically too long. Transmission problems were skirted by designing the memory so the wires were physically short and by keeping the inductance and capacitance on the lines as low as possible, thus minimizing their electrical length.

Second, the ferrite cores, if forced to switch back and forth too rapidly, can overheat, and heating seriously affects the magnetic characteristics of the ferrite material. Attaching the cores to an aluminum ground plane, which serves as a heat sink, overcame the heating problem.

The simplest layout for the bit plane—which con-

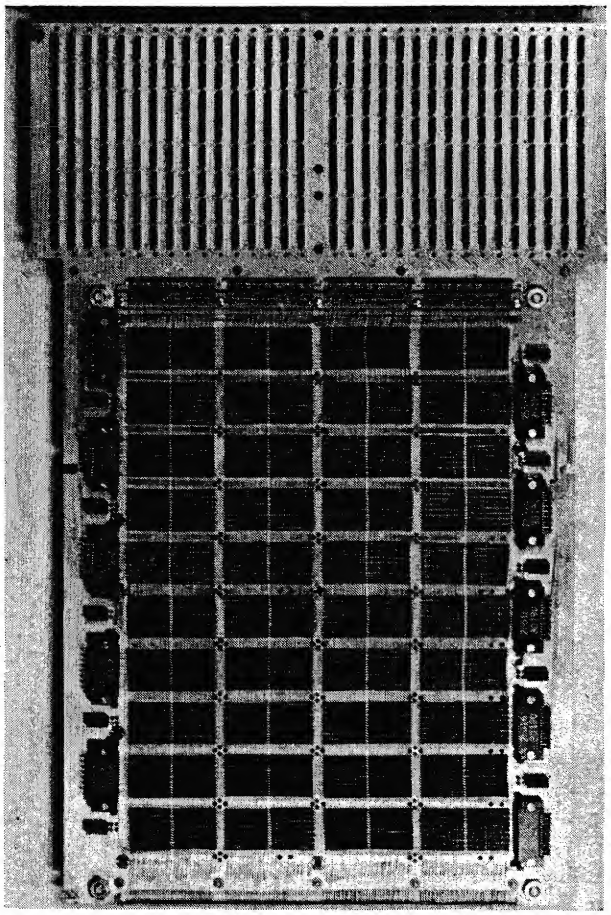

Core array for half of the 2½-D memory. Every small black square is a 32-by-32 group of cores with its own sense winding. Each group of eight small squares is one 32-by-256 bit plane, and has 16 bit windings making round trips through 32 rows. The entire memory comprises 20 bit planes on two ground plates, with 256 common word windings.

Reprinted from

Electronics | October 31, 1966

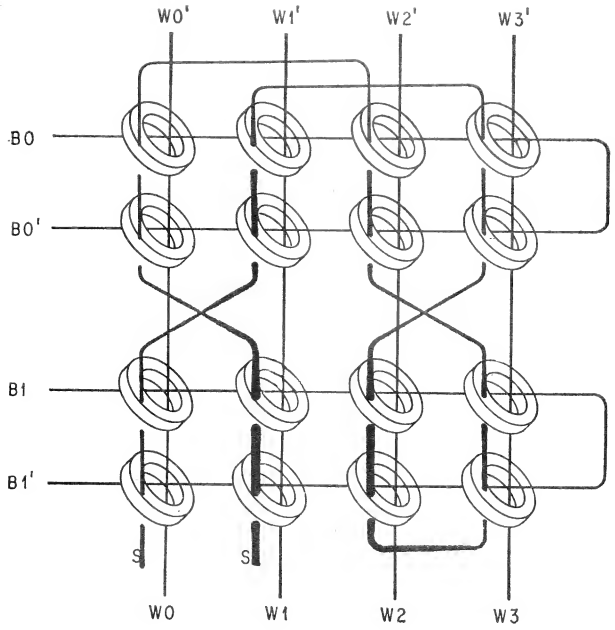

Sense winding run parallel to word windings, not bit windings as in some 2½-D memories. This diagram of a 4-by-4 array illustrates the winding layout and the "bow-tie" pattern. The inset shows how the bit line, lying between the sense and word lines, reduces coupling between them.

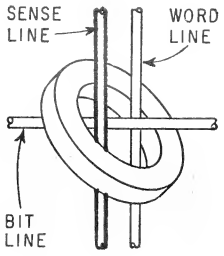

tains one bit of every word in the memory—is a square or nearly square array. But in a 2½-D memory the word winding must necessarily be much longer than the bit winding. The engineers reasoned that the bit plane could be made a long, narrow rectangle reducing the number of bit drive circuits and shortening the word winding considerably.

The engineers reduced inductance of the word and bit lines in two ways: they ran them close to the ground plane, thus putting the heat sink to double use; and they kept the bit line within the bit plane. The latter not only reduced loop inductance of the bit line, but eliminated the need for many solder connections required if the bit line was routed through more than one bit plane. And they reduced the capacitance of the word line by driving it from an unreferenced transformer instead of directly from the driving circuit.

The result was the new ferrite-core memory module, which can be combined into memory systems of almost unlimited word length and capacity.

One less wire

The 2½-D organization was selected for the new memory primarily for technical reasons; cost advantages were considered secondary. One technical advantage is the need for fewer wires through each core. In the 3-D organization, four wires pass through each core—the x-wire, the y-wire, the inhibit wire and the sense wire. Both the x-wire and the y-wire thread through their corresponding row and column in each bit plane. In the 2½-D organization, only three wires go through each core—the word wire, the bit wire and the sense wire. Only the word wire must pass through the corresponding columns in each bit plane. The bit wire terminates within the bit plane.

The inhibit wire limits the speed of the 3-D system in three major ways. First, the distributed capacitance causes degradation of the waveform, as described previously. Second, a voltage transition on the inhibit line inductively and capacitively couples unwanted signals into the sense line. Third, the current pulse on the inhibit line affects all the cores in the plane and induces delta noise in the sense winding. Delta noise is uncancelled noise from half-selected cores arising from the difference in slope between the top and bottom of the hysteresis loop; it is caused by selection currents as well as inhibit currents. Both the coupled-in noise and the delta noise impress unwanted signals on the sense winding, so that a significant interval of time must pass before the sense line recovers from the impressed noise and can again detect desired signals.

The inhibit pulse must rise before the coordinate drive pulse rises and fall after the latter falls, requiring additional cycle time. In the 2½-D memory organization, if a zero is to be written in the core, the current pulse on the bit drive lines is simply omitted during the write portion of the cycle. No bracketing or overlap of the pulses is necessary, reducing the cycle time.

An additional advantage realized by having only three wires through each core is that smaller cores can be used. Smaller cores require less drive current for a given switch time. In the 2½-D memory shorter wires threading fewer cores increase speed, because the propagation delays are shorter.

The single major advantage of 3-D organization is the economy in drive elements when compared to 2-D and 2½-D organization. This advantage, however, is less significant and may disappear entirely when attempts are made to increase both the speed and size of 3-D memories. Large 3-D memories need relatively long drive lines, requiring long signal propagation times. For high speed, the 3-D memory must be divided into smaller segments, each with its own complement of drive and sense circuits; and up goes the cost.

Less noise, less delay

The new memory's cycle time is short partly because arrangement of the core windings and their connections to driving circuits are designed to minimize the recovery time from noise coupled into the sense winding. Ideally, cycle time would be exactly twice the switching time of a single core—one switching action to fetch data and a second one to store new data or regenerate the old; in practice, the cycle time of any memory is longer because of address decoding, current rise times and the need to overcome noise coupling.

Address decoding and memory drive circuits were designed to reduce circuit delays as much as

Memory organizations

Computer memories can be organized in a number of different ways.

The most straightforward organization is called 2-dimensional or linear select. A 2-dimensional memory has all bits of one word in one plane. Two orthogonal sets of wires thread the plane, as shown at right; one set contains a wire for each word in the memory and the other set includes a wire for each bit in a word. At each intersection of the two sets, one word wire and one bit wire pass through the hole in a toroidal ferrite core. A third set of wires, the sense wires, is parallel to the set of bit wires and threads all the cores.

To fetch data from a 2-D memory, a full current (strong enough to switch to 0 all the cores that it threads) flows along the selected word wire. Cores that switch because they stored 1 bits generate voltage pulses in the sense wires. To store data, half the current passes in the reverse direction along the word wire and another half current passes along those bit wires threading cores that are to store 1 bits. The two currents combine to set the proper cores to the 1 state; where only the half current flows, the cores remain unswitched.

A 2-dimensional memory organization is expensive because for high speed a capacity of 2^n words requires 2^n drive elements.

A 3-dimensional, or coincident-current, memory consists of a series of stacked planes as shown at right; each plane contains one bit in each word in the memory. In each plane two orthogonal sets of wires permit any core in the plane to be addressed; the corresponding wires in all the planes are connected in series to select all the bits in a single word. In each plane a sense wire detects the change in magnetization of any core in that plane. An inhibit wire opposes the current in one of the two select wires when a zero is to be stored in the selected core in that plane.

To read data from a 3-D memory, a half current passes alone one wire in each of the two orthogonal sets, in such a direction as to switch to 0 all cores that they both thread. Again, cores that switch generate voltage pulses in the sense wires; where only the half current flows, the cores remain unswitched. To write data, a half current runs through each of the two wires in the reverse direction, switching cores to the 1 state except where the inhibit current opposes them. An inhibit pulse leaves a 0 stored in the core.

The 3-dimensional organization requires many fewer drive elements than the 2-D arrangement and is therefore less expensive; but it is also slower.

The 2½-dimensional organization like the 3-D, is a coincident-current memory, but it may resemble the 2-D more closely in its physical appearance. One possible arrangement is shown at right. Half currents pass through the word and bit wires in the proper direction to switch the desired core to the 0 state when reading or the 1 state when writing. To store a 0, the half current in the bit wire is omitted; no inhibit wire is needed. Sense wires thread the cores to detect one bit of every word stored in the memory; their exact arrangement varies in different designs.

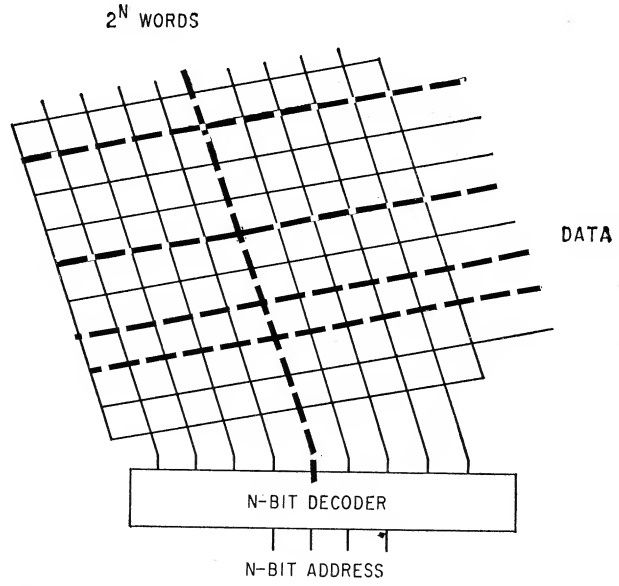

Two-dimensional memory plane.

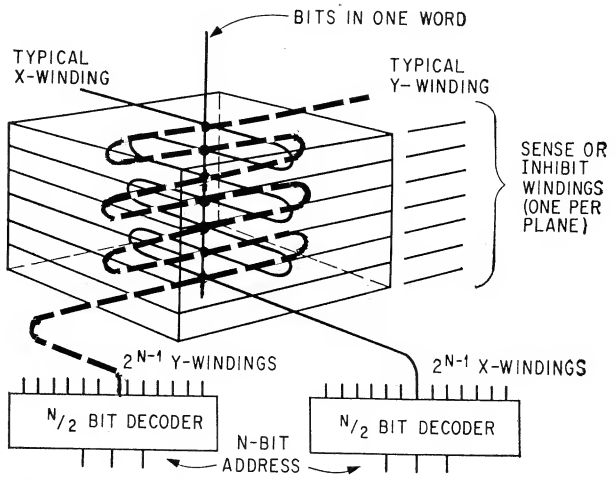

Three-dimensional memory, a stack of planes.

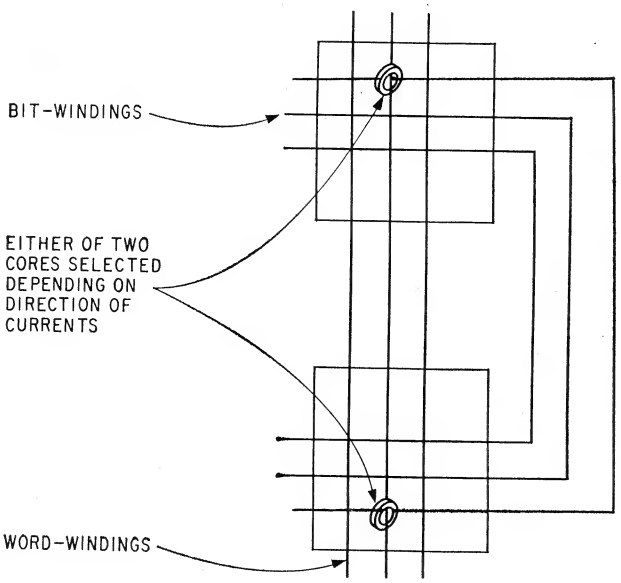

2½-dimensional memory combines 2-D and 3-D features.

the state of the semiconductor art permitted.

The rise and fall times of the current waveforms account for a significant portion of the over-all memory cycle time. In the new memory, inductance of the bit wire is kept down and the current rise time reduced by threading each bit wire through two immediately adjacent rows of cores, as shown on page 84. This minimizes the loop area enclosed by the wire and also provides a very short return path for the bit currents. The inductance is further reduced by keeping both wires close to the ground plane. Loop inductance is only 0.9 microhenry, as a result; without the ground plane it would be 1.7 μh.

The selected bit line and word line intersect at two cores because the bit wire loops back. The polarity of the bit current is always the same; the phasing of the word current relative to the bit current determines which of the two cores is selected. Looping back the bit line saves half the bit drive circuits.

Voltage and current transitions on the drive lines are capacitively and inductively coupled as noise into the sense lines. In some designs the rise of the word and bit currents is staggered to reduce the effects of this coupling and of delta noise, lengthening the cycle. In the new ½-microsecond memory, this staggering is not needed for three reasons: first, the sense line is perpendicular to the bit lines and parallel to the word lines, as shown on page 84. This greatly reduces the capacitive and inductive coupling between bit and sense lines. Second, capacitive coupling between the word and sense lines is significantly reduced by driving them through transformers whose secondary windings are not grounded or otherwise referenced; they are electrically floating. The transformer drive also cuts propagation delay by about half by providing a balanced drive. Third, the inductive coupling between the word and sense lines is minimized by placing the bit lines between them, providing extra physical separation (see inset, p. 84).

The amount of delta noise depends on the number of cores coupled by the sense line and its configuration. The more nearly square the core array for a given number of cores, the less delta noise is generated. For these reasons a square core array of 32 bits by 32 was selected for the memory design.

Columns and rows

The photograph on page 83 shows how the ½-microsecond memory stack is put together. Each of the small squares is one of the 32-by-32 array of cores; each array is threaded by a single sense winding. Eight arrays in a horizontal row form a 256 by 32 bit plane that contains 8,192 cores, corresponding to one bit in each of 8,192 words. Word wires enter the array on the 256 side; bit wires on the 32 side. The ten horizontal rows—mounted

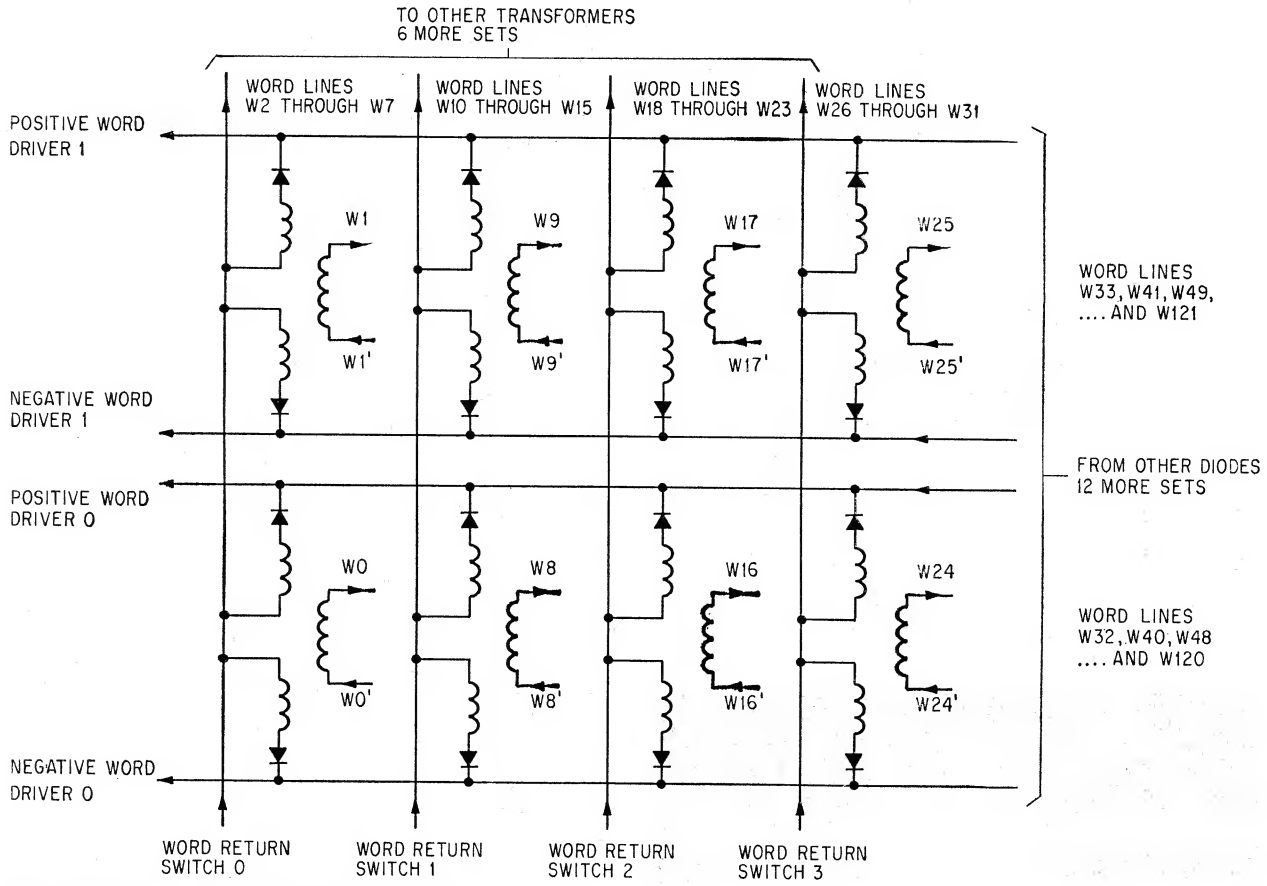

Unreferenced secondary windings of transformers generate the word current for the 2½-D memory. The floating secondaries reduce capacitive coupling to the sense winding and provide a balanced drive to the word winding.

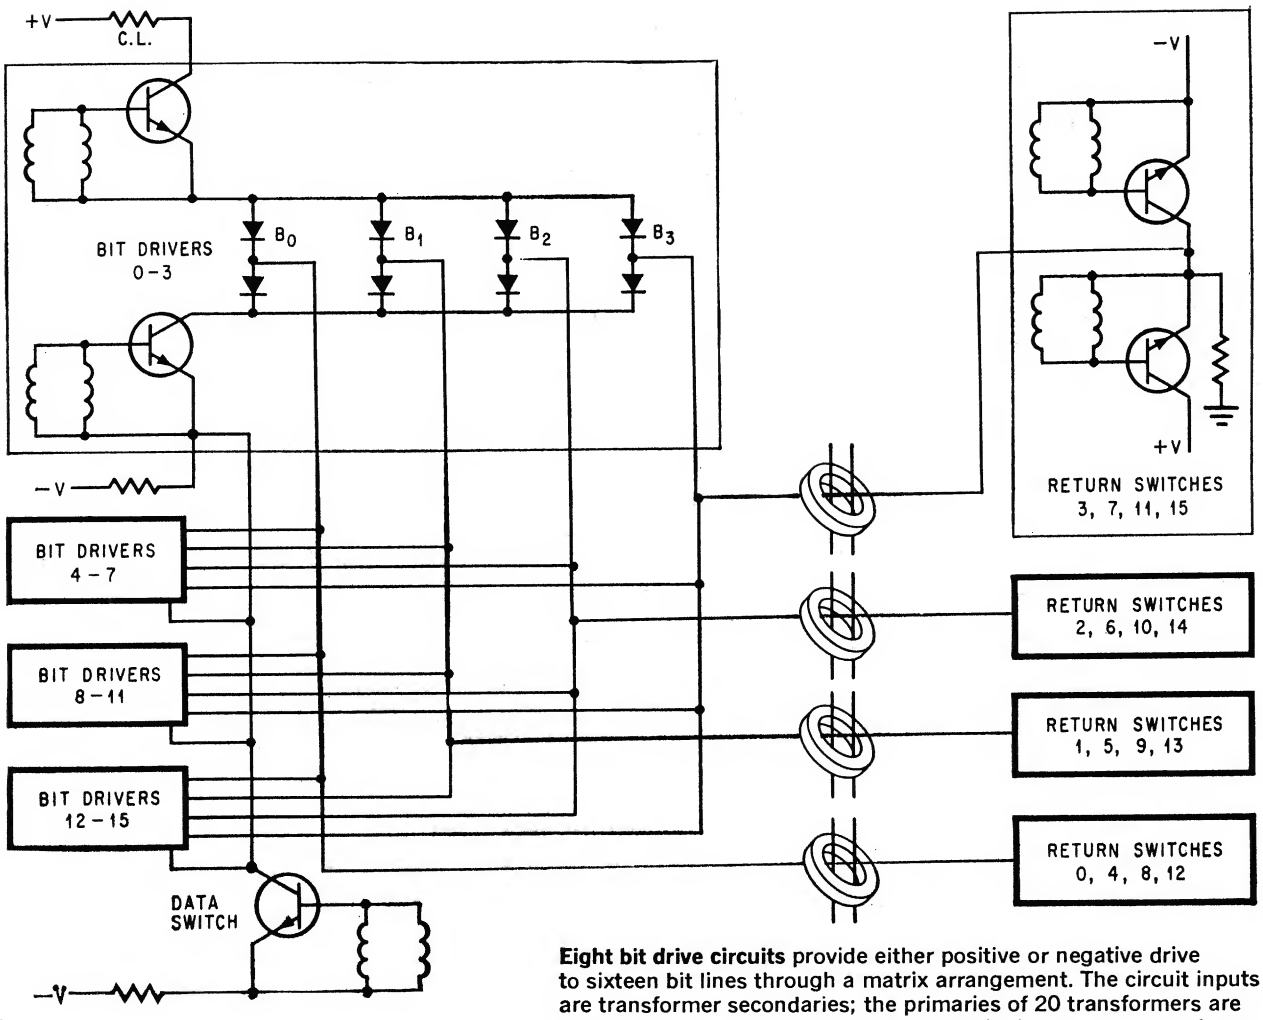

Eight bit drive circuits provide either positive or negative drive to sixteen bit lines through a matrix arrangement. The circuit inputs are transformer secondaries; the primaries of 20 transformers are connected in series to insure that all 20 bits in a memory word are fetched or stored simultaneously. The inductance shown across the secondary winding turns off the transistor quickly.

on a single base plate—contain ten bits in all the words. Two base plates together make up the complete ½-microsecond memory of 8,192 twenty-bit words; the 256 word wires continue through the two plates.

The cores in the memory are of 0.020-inch outside diameter and 0.012-inch inside diameter, placed 0.025 inch apart center to center. They require a nominal 840-milliamperes drive current and switch in 120 nanoseconds. Their output waveform is shown below.

Word and bit drivers

The ungrounded secondary winding of a transformer drives each of the 256 word lines in the memory, reducing capacitive coupling to the sense lines. Each transformer has two primary windings—one provides current for reading and the other for writing; the two currents are of opposite polarity. The transformers are arranged in two conventional matrixes of 8 by 16 transformers each, as shown on the opposite page; two separate matrixes present a smaller loop inductance to the word drive circuits than one large matrix. One constant-current driver and one switch turn on to read or write one word. Propagation delay in the word line is less than four nanoseconds.

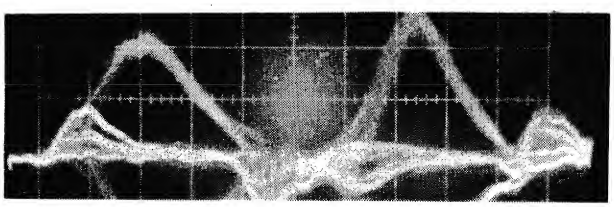

Output waveforms reading a continuous series of 1's. The scale is 40 nanoseconds per division.

Bit drive circuits are shown in the diagram above. Each group of 16 bit lines is driven by its own 16-way switch matrix. Baluns are used in the drive circuit, one balun for each bit line to limit power dissipation of the driver switches. The 20-bit memory contains 20 of these matrixes. Each group of 16 lines doubles back through adjacent rows of cores to provide the 32 bit lines in the 256-by-32 core group. The inputs to the drivers are another set of transformer secondaries; the transformer primaries in corresponding drive and switch circuits are connected in series to insure that circuits turn on simultaneously and minimize the

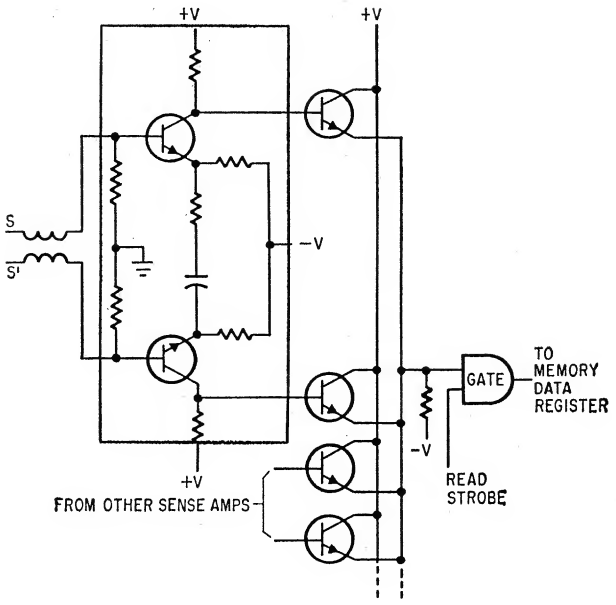

Sense amplifier connects eight sense windings to a single strobe gate through emitter-followers whose outputs share a common load. Differential preamplifiers (one of which is outlined in color) drive the emitter followers. This configuration keeps down the cost of sensing the many small sense sections in the memory.

number of decoding circuits. Corresponding drive and switch circuits turn on in all 20 matrixes to read or write one word. The current-limiting resistors (marked C.L. in the diagram) determine the current in each line. During the cycle's write portion, an independent data switch in each bit matrix controls whether a 1 or a 0 is written into the selected bit; if the data switch is not turned on, no current is available for the bit line. The propagation delay in the bit line is two nanoseconds.

Sensing data

Each sense wire is threaded through only a 32-by-32 array to make the 500-nanosecond cycle time possible. The small sense section requires eight sense preamplifiers per bit, or a total of 160 sense preamplifiers in the memory; their cost must therefore be low to keep the cost of the entire memory from getting out of hand.

A schematic of one sense amplifier is shown above, with the connections of seven other preamplifiers indicated to produce a single bit output. The input passes through a transformer connected as a balun to eliminate common-mode noise. The differential stage, outlined in color, is a hybrid circuit made of discrete transistors and screen-printed thick-film resistors on a ceramic substrate. The base-to-emitter voltage and other parameters vary from transistor to transistor; to compensate for these critical variations, the collector operating point and small-signal gain are adjusted by trimming the resistors.

Following the differential stage is a pair of emitter followers with common emitters and a single load resistor. This configuration combines the outputs of the differential stage's two sides as in an OR gate. Seven other emitter-follower pairs are connected to the same common load, bringing the outputs of the eight sense windings and eight preamplifiers to a single common point. Only one of the sixteen emitters will be active at any one time, because only one core in each bit plane switches at one time. A strobe gate samples the signal at the common point to reduce still further the likelihood of false outputs generated by noise. The strobed and clipped signal is then made available to the memory data register.

Packaging for speed

Memory speed depends to a great extent on the packaging arrangement, because packaging affects line lengths, grounding and voltage distribution, among other things.

The entire memory is mounted on two aluminum base plates that serve both as a heat sink and a ground plane. The cores are secured to the plate by a silicon and magnesium-oxide mixture that transfers heat from the cores to the plate.

Bit selection diodes and baluns are mounted on printed-circuit boards that plug into the connectors along the side of the core array. In the photo on page 83, part numbers are visible on the sides of these connectors; in place the cards project toward the array's center, covering up the cores. The bit drive circuits are placed on other printed-circuit boards that lie flat on top of the cards carrying the diodes and baluns; connectors establish contact between the boards. The word drive circuits are on p-c boards that plug into the connector array at one end of the core array. Boards carrying the sense amplifiers are connected in the frame that holds the rest of the assembly; these connectors are not visible in the photo. This assembly provides the shortest possible interconnections.

Voltages are distributed on strip-lines fabricated as part of the printed-circuit cards that carry drive and sense circuitry. The strip-lines reduce the liability of noise spikes and the need for filtering. Critical time pulses—the read strobe signal, for instance—are also distributed on strip-line. This line has a propagation time similar to that of the word current through the core array, so that the strobe pulse and the memory output remain fixed in time relative to each other throughout the entire memory.

The entire memory measures 26½ by 19¼ inches, and the complete assembly is 4¼ inches thick. These measurements include the core arrays, the ground planes on which they are mounted, the surrounding drive and sense circuitry, and the supporting frame.

The author

Alexander Elovic is the manager of memory systems at the Electronic Components division of the Burroughs Corp. He has had the job for about a year, having previously been the director of engineering at Indiana General Corp.